Bayesian Exponential Random Graph Models with Nodal Random Effects

S. Thiemichen[⋆], N. Friel[†], A. Caimo[‡], G. Kauermann[⋆].

[⋆]*Institut für Statistik Ludwigs-Maximilians-Universität München, Germany.*

[†]*School of Mathematical Sciences and Insight: The National Centre for Data Analytics, University College Dublin, Ireland.*

[‡]*Social Network Analysis Research Center, Faculty of Economics, University of Lugano, Switzerland.*

January 13, 2015

arXiv:1407.6895v2 [stat.AP] 12 Jan 2015

Abstract

We extend the well-known and widely used Exponential Random Graph Model (ERGM) by including nodal random effects to compensate for heterogeneity in the nodes of a network. The Bayesian framework for ERGMs proposed by Caimo and Friel (2011) yields the basis of our modelling algorithm. A central question in network models is the question of model selection and following the Bayesian paradigm we focus on estimating Bayes factors. To do so we develop an approximate but feasible calculation of the Bayes factor which allows one to pursue model selection. Two data examples and a small simulation study illustrate our mixed model approach and the corresponding model selection.

1 Introduction

The analysis of network data is an emerging field in statistics which is challenging both model-wise and computationally. Recently Goldenberg et al. (2010), Hunter et al. (2012), Fienberg (2012), and Salter-Townshend et al. (2012), respectively, published comprehensive survey articles discussing statistical approaches, challenges and developments in network data analysis. We also refer to the monograph of Kolaczyk (2009) for a comprehensive introduction to the field.

In this paper we consider networks represented as a $n \times n$ dimensional adjacency matrix $\boldsymbol{Y}$, where the element $Y_{ij} = 1$, if an edge exists between vertex i and vertex j and $Y_{ij} = 0$ otherwise, with $i, j \in \{1, \ldots, n\}$ and $i \neq j$, that is there is no connection from a vertex to itself. With n we denote the number of vertices in the network and for simplicity we assume undirected edges, that is $Y_{ij} = Y_{ji}$. Therefore, the matrix $\boldsymbol{Y}$ is symmetric and for simplicity it is sufficient to consider the upper triangle of $\boldsymbol{Y}$ only, that is $Y_{ij}, j > i$. Our approach equally applies to non-symmetric adjacency matrices corresponding to directed graphs. A concrete realisation of $\boldsymbol{Y}$ is denoted with $\boldsymbol{y}$.

With respect to the available statistical models for modelling cross-sectional network data one may roughly distinguish between two strands, (a) models which explain the existence

of an edge purely with external nodal covariates or random effects and (b) models where the existence of an edge also depends on the local network structure. The first strand of models is phrased as p_1 and p_2 models tracing back to Holland and Leinhardt (1981). Specifically, in the p_1 model we set

$$\text{logit}\left[\mathbb{P}(Y_{ij}=1)\right] = \log\left\{\frac{\mathbb{P}(Y_{ij}=1)}{1-\mathbb{P}(Y_{ij}=1)}\right\} = \alpha_i + \alpha_j + \boldsymbol{z}_{ij}^t\boldsymbol{\beta} \tag{1}$$

where $\boldsymbol{z}_{ij}$ denotes a set of covariates relating to the vertices i and j and α_i and α_j are nodal effects, here assuming undirected edges. Since the number of parameters increases with increasing network size n, Duijn et al. (2004) proposed to replace the α parameters in (1) by random effects, see also Zijlstra et al. (2006). This yields the p_2 model

$$\begin{aligned}\text{logit}\left[\mathbb{P}(Y_{ij}=1|\boldsymbol{\phi})\right] &= \phi_i + \phi_j + \boldsymbol{z}_{ij}^t\beta, \\ \boldsymbol{\phi} = (\phi_1,\ldots,\phi_n)^t &\sim N(0,\sigma_\phi^2 I_n)\end{aligned} \tag{2}$$

with I_n as n dimensional unit matrix. A general principle with this approach is that vertices (or actors in the network, respectively) are not considered as homogeneous but heterogeneous, though their heterogeneity is not observable but latent and expressed in the node specific random effects $\phi_i,\ i=1,\ldots,n$.

Both, the p_1 and the p_2 model lie within the classical generalized linear (mixed) model framework which allows estimation using standard statistical software. The p_2 models also allow for Bayesian estimation approaches, see for example Gill and Swartz (2004).

The second strand in statistical network modelling is based on the so called Exponential Random Graph Model (ERGM) proposed by Frank and Strauss (1986). Here we model directly the network using the likelihood function

$$\mathbb{P}(\boldsymbol{Y}=\boldsymbol{y}|\boldsymbol{\theta}) = f(\boldsymbol{y}|\boldsymbol{\theta}) = \frac{q_{\boldsymbol{\theta}}(\boldsymbol{y})}{\kappa(\boldsymbol{\theta})} = \frac{\exp\left\{\boldsymbol{\theta}^t s(\boldsymbol{y})\right\}}{\kappa(\boldsymbol{\theta})} \tag{3}$$

where $\boldsymbol{\theta} = (\theta_1,\ldots,\theta_p)^t$ is the vector of model parameters and $s(\boldsymbol{y})$ is a vector of sufficient network statistics like the number of edges or two-stars in a network, see for example Snijders et al. (2006). In equation (3) the term $\kappa(\boldsymbol{\theta})$ denotes the normalizing constant, that is

$$\kappa(\boldsymbol{\theta}) = \sum_{\boldsymbol{y}\in\mathcal{Y}} \exp\left\{\boldsymbol{\theta}^t s(\boldsymbol{y})\right\}$$

and is accordingly the sum over $2^{\binom{n}{2}}$ potential undirected graphs and therefore numerically intractable, except for very small graphs. Early fitting approaches are based on the pseudolikelihood idea proposed by Strauss and Ikeda (1990). More advanced are MCMC based routines proposed by Hunter and Handcock (2006) based on the work of Geyer and Thompson (1992). A fully Bayesian approach to estimate ERGMs has been developed by Caimo and Friel (2011).

Model (3) allows for a conditional interpretation by focusing on the occurrence of a single edge between two nodes. To be specific we obtain

$$\text{logit}\left[\mathbb{P}\Big(Y_{ij}=1|Y_{kl},(k,l)\neq(i,j);\boldsymbol{\theta}\Big)\right] = \boldsymbol{\theta}^t\ s_{ij}(\boldsymbol{y}), \tag{4}$$

where $s_{ij}(\boldsymbol{y})$ denotes the vector of so called change statistics

$$s_{ij}(\boldsymbol{y}) = s\Big(y_{ij} = 1, y_{kl}, (k,l) \neq (i,j)\Big) - s\Big(y_{ij} = 0, y_{kl}, (k,l) \neq (i,j)\Big).$$

We refer to Robins et al. (2007a) and Robins et al. (2007b) for a deeper discussion of Exponential Random Graph Models.

Contrasting equation (4) with the p_1 and p_2 model given in equations (1) and (2) it becomes obvious that the ERGM in contrast to the p_1 and p_2 models take the network structure into account while considering the nodes to be homogeneous. When modelling network data this means that all possible heterogeneity in the network nodes (that is the actors in the network) is included as covariates in the model and influence the (global) structure of the network. Since homogeneity of the nodes have led from p_1 to p_2 models, we want to pursue the same modelling exercise by allowing for latent node specific heterogeneity in Exponential Random Graph Models. To do so, we combine the p_2 model (2) with the ERGM (4) towards

$$\text{logit}\left[\mathbb{P}\Big(Y_{ij} = 1 | Y_{kl}, (k,l) \neq (i,j); \boldsymbol{\theta}, \phi_i, \phi_j\Big)\right] = \boldsymbol{\theta}^t \; s_{ij}(\boldsymbol{y}) + \phi_i + \phi_j \tag{5}$$

with $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_n)^t$ and $\phi_i \overset{\text{i.i.d.}}{\sim} N(\mu_\phi, \sigma_\phi^2),\ i = 1, \ldots, n$. The parameter μ_ϕ captures the average propensity in the network for forming a tie. In terms of the likelihood function for the whole network we obtain from (5)

$$\mathbb{P}(\boldsymbol{Y} = \boldsymbol{y} | \boldsymbol{\theta}, \boldsymbol{\phi}) = f(\boldsymbol{y} | \boldsymbol{\theta}, \boldsymbol{\phi}) = \frac{q_{\boldsymbol{\theta},\boldsymbol{\phi}}(\boldsymbol{y})}{\kappa(\boldsymbol{\theta}, \boldsymbol{\phi})} = \frac{\exp\left\{\boldsymbol{\theta}^t s(\boldsymbol{y}) + \boldsymbol{\phi}^t t(\boldsymbol{y})\right\}}{\kappa(\boldsymbol{\theta}, \boldsymbol{\phi})}, \tag{6}$$

where $t(\boldsymbol{y})$ contains the degree statistics of the n vertices, i.e. $t_i(\boldsymbol{y}) = \sum_{j=1}^{n} y_{ij},\ \text{for } i = 1, \ldots, n$. That is we fit an Exponential Random Graph Model with random, node specific effects accounting for heterogeneity. The model in equations (5) and (6) falls in the general class of Exponential-family Random Network Models proposed by Fellows and Handcock (2012) but unlike their model we treat the node specific effect as latent and we pursue a fully Bayesian estimation. We also refer to Krivitsky et al. (2009) who propose a model with actor specific random effects based on a latent cluster model. The authors also propose node specific random effects. We follow this line and give further interpretability of the effects. A central issue in model extensions is the question of model selection. We emphasize this point in the paper by comparing models with and without nodal effects using the Bayes factor as model selection criterion. However, calculation of the Bayes factor suffers from the above mentioned problem in Exponential Random Graph Models in that the normalization constant $\kappa(\cdot)$ is numerically infeasible. We therefore propose an approximate calculation of the Bayes factor and show in a simulation study its usability for model selection.

For estimation and model selection of model (6) we extend the fully Bayesian approach from Caimo and Friel (2011). The developed estimation routine is based on the numerical work of Caimo and Friel (2014) with their R (R Core Team, 2014) package `Bergm` (see http://cran.r-project.org/web/packages/Bergm). Our algorithms for model fitting and

selection will be included in the `Bergm` package.

The paper is organized as follows. In Section 2 we derive a fully Bayesian formulation of the model. This is followed by a detailed description of the MCMC based estimation routine. Section 3 deals with the issue of model selection using Bayes factors. Two data examples and some simulation results are presented Section 4. Finally Section 5 concludes with a discussion.

2 Bayesian model formulation and Estimation

Before proposing a fully Bayesian formulation for model (6) bear in mind that the normalizing constant $\kappa(\boldsymbol{\theta}, \boldsymbol{\phi})$ is numerically infeasible to calculate except for small networks so that numerically demanding simulation based fitting routines need to be employed. We follow a fully Bayesian approach by imposing a prior distribution on $\boldsymbol{\theta}$. The posterior of interest for the Bayesian Exponential Random Graph Model with nodal random effects in (6) then becomes

$$p(\boldsymbol{\theta}, \boldsymbol{\phi}, \mu_\phi, \sigma^2_\phi | \boldsymbol{y}) = \frac{f(\boldsymbol{y}|\boldsymbol{\theta}, \boldsymbol{\phi}) p(\boldsymbol{\theta}) p(\boldsymbol{\phi}|\mu_\phi, \sigma^2_\phi) p(\mu_\phi) p(\sigma^2_\phi)}{p(\boldsymbol{y})}, \tag{7}$$

where $p(\boldsymbol{\theta})$ is the prior distribution of $\boldsymbol{\theta}$ and $p(\boldsymbol{\phi}|\mu_\phi, \sigma^2_\phi)$ the prior for the random nodal effects $\boldsymbol{\phi}$. We assume the nodal effects to be independent and identically normally distributed, that is

$$\phi_i \sim N(\mu_\phi, \sigma^2_\phi), \quad \text{for } i = 1, \ldots, n$$

and accordingly we use $\boldsymbol{\theta} \sim N(0, \rho^2 I_p)$, with I_p denoting the p-dimensional unity matrix and ρ^2 chosen such that the prior distribution is flat. For the hyper prior distribution $p(\mu_\phi)$ of the mean μ_ϕ we assume a normal distribution centred at 0, that is

$$\mu_\phi \sim N(0, \tau^2).$$

The hyper prior $p(\sigma^2_\phi)$ of the variance σ^2_ϕ is assumed to be an inverse gamma distribution, that is

$$\sigma^2_\phi \sim IG(a, b).$$

Finally, the parameters τ^2, a and b are all constants and chosen in a way that results in flat hyper prior distributions. Figure 1 illustrates this Bayesian model formulation.

It is important to note, that the posterior distribution in (7) is so-called doubly-intractable. This is because, firstly, it is not possible to evaluate the posterior density (7) due to $p(\boldsymbol{y})$, the marginal likelihood or evidence, being intractable. Secondly, it is also numerically infeasible to calculate the normalizing constant $\kappa(\boldsymbol{\theta}, \boldsymbol{\phi})$ in the likelihood $f(\boldsymbol{y}|\boldsymbol{\theta}, \boldsymbol{\phi})$ except for very small network graphs. Similar to the algorithm proposed by Caimo and Friel (2011) we use the so-called exchange algorithm from Murray et al. (2006) to draw samples from the posterior distribution of interest. Let therefore $\boldsymbol{\gamma} = (\boldsymbol{\theta}, \boldsymbol{\phi})$ denote the entire parameter vector of the ERGM. Instead of drawing directly from (7),

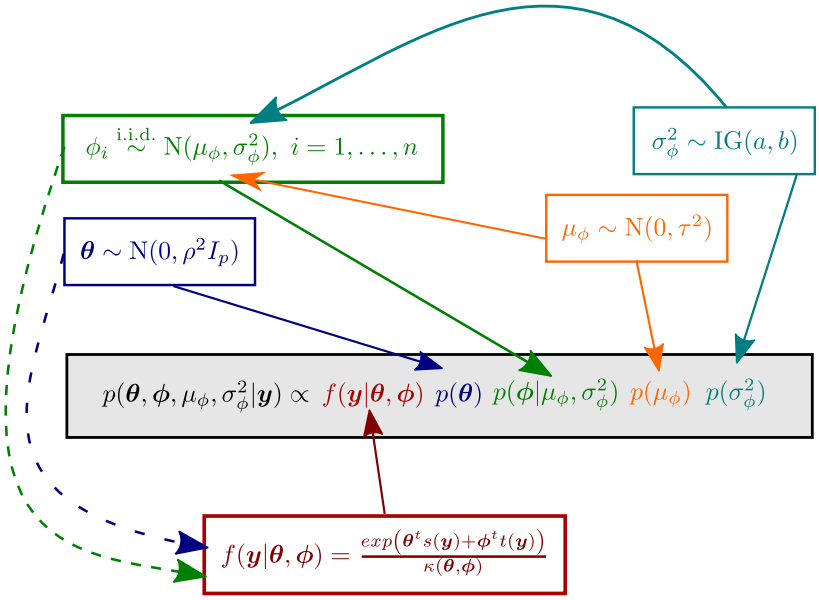
$\phi_i \overset{\text{i.i.d.}}{\sim} \mathrm{N}(\mu_\phi, \sigma_\phi^2),\ i = 1, \ldots, n$
$\sigma_\phi^2 \sim \mathrm{IG}(a, b)$
$\mu_\phi \sim \mathrm{N}(0, \tau^2)$
$\boldsymbol{\theta} \sim \mathrm{N}(0, \rho^2 I_p)$
$p(\boldsymbol{\theta}, \boldsymbol{\phi}, \mu_\phi, \sigma_\phi^2 | \boldsymbol{y}) \propto f(\boldsymbol{y}|\boldsymbol{\theta}, \boldsymbol{\phi})\ p(\boldsymbol{\theta})\ p(\boldsymbol{\phi}|\mu_\phi, \sigma_\phi^2)\ p(\mu_\phi)\ p(\sigma_\phi^2)$
$f(\boldsymbol{y}|\boldsymbol{\theta}, \boldsymbol{\phi}) = \frac{exp(\boldsymbol{\theta}^t s(\boldsymbol{y}) + \boldsymbol{\phi}^t t(\boldsymbol{y}))}{\kappa(\boldsymbol{\theta}, \boldsymbol{\phi})}$

Figure 1 Overview of the Bayesian model formulation for the Exponential Random Graph Model with nodal random effects.

we sample from the augmented distribution

$$\begin{aligned} &p(\boldsymbol{\gamma}', \boldsymbol{y}', \boldsymbol{\gamma}, \mu_\phi, \sigma_\phi^2 | \boldsymbol{y}) \propto \\ &\quad f(\boldsymbol{y}|\boldsymbol{\gamma}) p(\boldsymbol{\gamma}|\mu_\phi, \sigma_\phi^2) p(\mu_\phi) p(\sigma_\phi^2) h(\boldsymbol{\gamma}'|\boldsymbol{\gamma}) f(\boldsymbol{y}'|\boldsymbol{\gamma}'), \end{aligned} \tag{8}$$

where $h(\cdot|\cdot)$ is a proposal function, to be specified later. This proposal provides $\boldsymbol{\gamma}' = (\boldsymbol{\theta}', \boldsymbol{\phi}')$ as new candidate values for $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$, respectively, and based on $\boldsymbol{\gamma}'$ we can simulate $\boldsymbol{y}'$ as an auxiliary network. The proposal is accepted with probability

$$\alpha = \min\left(1, \frac{q_\gamma(\boldsymbol{y}') p(\boldsymbol{\gamma}') h(\boldsymbol{\gamma}|\boldsymbol{\gamma}') q_{\gamma'}(\boldsymbol{y})}{q_\gamma(\boldsymbol{y}) p(\boldsymbol{\gamma}) h(\boldsymbol{\gamma}'|\boldsymbol{\gamma}) q_{\gamma'}(\boldsymbol{y}')} \times \frac{\kappa(\boldsymbol{\gamma})\kappa(\boldsymbol{\gamma}')}{\kappa(\boldsymbol{\gamma})\kappa(\boldsymbol{\gamma}')}\right) \tag{9}$$

where $p(\boldsymbol{\gamma}) = p(\boldsymbol{\theta}) \cdot p(\boldsymbol{\phi}|\mu_\phi, \sigma_\phi^2)$. Note that in (9) the normalizing constants cancel out so that (9) is in principle easy to calculate. Though the algorithm is in this form a direct extension of the BERGM algorithm in Caimo and Friel (2011) it is advisable to separate the proposals of $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ to achieve higher acceptance rates. This is described in the following algorithmic steps. In detail, our algorithm works as follows:

Algorithm 1: Fit BERGM with nodal random effects

Step 1: Gibbs update of $(\boldsymbol{\theta}', \boldsymbol{y}')$:

(i) Draw $\boldsymbol{\theta}' \sim h(\cdot|\boldsymbol{\theta})$.

(ii) Draw $\boldsymbol{y}' \sim p(\cdot|\boldsymbol{\theta}', \boldsymbol{\phi})$.

(iii) Propose to move from $\boldsymbol{\theta}$ to $\boldsymbol{\theta}'$ with probability

$$\alpha = \min\left(1, \frac{q_{\boldsymbol{\theta},\boldsymbol{\phi}}(\boldsymbol{y}')p(\boldsymbol{\theta}')h(\boldsymbol{\theta}|\boldsymbol{\theta}')q_{\boldsymbol{\theta}',\boldsymbol{\phi}}(\boldsymbol{y})}{q_{\boldsymbol{\theta},\boldsymbol{\phi}}(\boldsymbol{y})p(\boldsymbol{\theta})h(\boldsymbol{\theta}'|\boldsymbol{\theta})q_{\boldsymbol{\theta}',\boldsymbol{\phi}}(\boldsymbol{y}')} \times \frac{\kappa(\boldsymbol{\theta},\boldsymbol{\phi})\kappa(\boldsymbol{\theta}',\boldsymbol{\phi})}{\kappa(\boldsymbol{\theta},\boldsymbol{\phi})\kappa(\boldsymbol{\theta}',\boldsymbol{\phi})}\right).$$

Step 2: Gibbs update of $(\boldsymbol{\phi}', \boldsymbol{y}')$:

(i) Draw $\boldsymbol{\phi}' \sim g(\cdot|\boldsymbol{\phi})$.

(ii) Draw $\boldsymbol{y}' \sim p(\cdot|\boldsymbol{\theta}, \boldsymbol{\phi}')$.

(iii) Propose to move from $\boldsymbol{\phi}$ to $\boldsymbol{\phi}'$ with probability

$$\alpha = \min\left(1, \frac{q_{\boldsymbol{\theta},\boldsymbol{\phi}}(\boldsymbol{y}')p(\boldsymbol{\phi}'|\mu_\phi, \sigma_\phi^2)g(\boldsymbol{\phi}|\boldsymbol{\phi}')q_{\boldsymbol{\theta},\boldsymbol{\phi}'}(\boldsymbol{y})}{q_{\boldsymbol{\theta},\boldsymbol{\phi}}(\boldsymbol{y})p(\boldsymbol{\phi}|\mu_\phi, \sigma_\phi^2)g(\boldsymbol{\phi}'|\boldsymbol{\phi})q_{\boldsymbol{\theta},\boldsymbol{\phi}'}(\boldsymbol{y}')} \times \frac{\kappa(\boldsymbol{\theta},\boldsymbol{\phi})\kappa(\boldsymbol{\theta},\boldsymbol{\phi}')}{\kappa(\boldsymbol{\theta},\boldsymbol{\phi})\kappa(\boldsymbol{\theta},\boldsymbol{\phi}')}\right).$$

Step 3: Metropolis-Hastings update of μ_ϕ:
Draw proposal μ'_ϕ from $k(\cdot|\mu_\phi)$ and accept the proposed value with probability $\alpha = \min\left(1, \frac{p(\boldsymbol{\phi}|\mu_\phi,\sigma_\phi^2)p(\mu_\phi)}{p(\boldsymbol{\phi}|\mu'_\phi,\sigma_\phi^2)p(\mu'_\phi)}\right)$.

Step 4: Metropolis-Hastings update of σ_ϕ^2:
Draw proposal $\sigma_\phi^{2\prime}$ from $l(\cdot|\sigma_\phi^2)$ and accept the proposed value with probability $\alpha = \min\left(1, \frac{p(\boldsymbol{\phi}|\mu_\phi,\sigma_\phi^2)p(\sigma_\phi^2)}{p(\boldsymbol{\phi}|\mu_\phi,\sigma_\phi^{2\prime})p(\sigma_\phi^{2\prime})}\right)$.

Start again with *Step 1* until the maximum number of iterations is reached.

It is easy to see that there is no necessity to compute the normalizing constants $\kappa(\cdot)$, because they cancel out when calculating the acceptance probabilities in the first two steps of the algorithm. The current implementation of the algorithm uses single-site updates for the update of $\boldsymbol{\phi}$, that is each $\phi_i,\ i = 1, \ldots, n$ is updated in turn while all other values are kept constant. This leads to reasonable acceptance probabilities for the Markov chain.
The default choices for the proposal functions $h(\cdot|\cdot)$, $g(\cdot|\cdot)$ and $k(\cdot|\cdot)$ are normal distributions centred at the current parameter value, for $l(\cdot|\cdot)$ we use a uniform distribution, which is symmetric around the current value of σ_ϕ^2 and truncated at zero to avoid negative proposals for the variance parameter.
The draws of the auxiliary network $\boldsymbol{y}'$ in the component of steps 1 and 3 are realised using the "tie no tie" sampler from the `ergm` package (Hunter et al., 2008), which is a simple Gibbs sampler. Although this auxiliary Gibbs sampler does not yield an exact draw $\boldsymbol{y}'$, Everitt (2012) has shown, under some assumptions, that the resulting approximate exchange algorithm converges to the target distribution as the number of auxiliary draws tends to infinity. As a practical result he points out that for the number of auxiliary iterations it is often sufficient to use roughly the number of possible ties in the network.

3 Model Selection

Model Selection is an important, often neglected issue in network data analysis. We put special emphasis on this task here and propose the Bayes factor suitable for model selection. One of the interesting questions in our model is, if we are able to distinguish the three following model generating processes:

(1) Nodal random effects only, i.e. the p_2 model,

(2) Structural effects only, i.e. the standard ERGM, and

(3) ERGM in combination with nodal random effects.

This question results in the problem of model selection. The data examples in Section 4.1 illustrate this issue.
Classical Bayesian tools for model comparison such as the deviance information criterion (DIC) as suggested by Spiegelhalter et al. (2002) are not directly available, again due to the intractability of the normalizing constant of the likelihood in model equation (6).
Computing Bayes factors for model choice using reversible jump Markov Chain Monte Carlo for Bayesian Exponential Random Graph Models as done by Caimo and Friel (2013) is not an option for our model. This approach would be possible in general, but very time consuming from a computational point of view.

We suggest the following strategy for deciding whether to include nodal random effects into the model or not. The goal is to calculate a Bayes factor for two competing models (Kass and Raftery, 1995). First we fit the two Exponential Random Graph Models, one with edges and non-random effects only, notated as model m_1 with coefficients $\boldsymbol{\theta}'$, and the second one with nodal random effects instead of the edges term, labelled as model m_2 with coefficients $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$. Note that these two models are nested.[1]
Following Bayes theorem the so-called evidence for each model can be calculated using

$$p(\boldsymbol{y}|m_1) = \frac{f(\boldsymbol{y}|\boldsymbol{\theta}')p(\boldsymbol{\theta}')}{p(\boldsymbol{\theta}'|\boldsymbol{y})}, \quad \forall\ \boldsymbol{\theta}', \tag{10}$$

for model m_1, and

$$\begin{aligned} p(\boldsymbol{y}|m_2) &= \frac{f(\boldsymbol{y}|\boldsymbol{\theta},\boldsymbol{\phi},\mu_\phi,\sigma^2_\phi)p(\boldsymbol{\theta})p(\boldsymbol{\phi}|\mu_\phi,\sigma^2_\phi)p(\mu_\phi)p(\sigma^2_\phi)}{p(\boldsymbol{\theta},\boldsymbol{\phi},\mu_\phi,\sigma^2_\phi|\boldsymbol{y})}, \quad \forall\ \boldsymbol{\theta},\boldsymbol{\phi},\mu_\phi,\sigma^2_\phi, \\ &= \frac{f(\boldsymbol{y}|\boldsymbol{\theta},\mu_\phi,\sigma^2_\phi)p(\boldsymbol{\theta})p(\mu_\phi)p(\sigma^2_\phi)}{p(\boldsymbol{\theta},\mu_\phi,\sigma^2_\phi|\boldsymbol{y})}, \quad \forall\ \boldsymbol{\theta},\mu_\phi,\sigma^2_\phi \end{aligned} \tag{11}$$

for model m_2.
The term $f(\boldsymbol{y}|\boldsymbol{\theta},\mu_\phi,\sigma^2_\phi)$ denotes the marginal likelihood from model m_2, where the random

[1] For the approach presented here, especially for the path sampling, we need the models to be nested. In general it would be possible to extend the approach to non-nested models as well.

effects $\boldsymbol{\phi}$ have been marginalized, i.e.

$$f(\boldsymbol{y}|\boldsymbol{\theta}, \mu_\phi, \sigma^2_\phi) = \int \frac{\exp\left\{\boldsymbol{\theta}^t s(\boldsymbol{y}) + \boldsymbol{\phi}^t t(\boldsymbol{y})\right\}}{\kappa(\boldsymbol{\theta}, \boldsymbol{\phi})} \cdot p(\boldsymbol{\phi}|\mu_\phi, \sigma^2_\phi)\ \mathrm{d}\boldsymbol{\phi}$$
$$\approx \frac{\exp\left\{\boldsymbol{\theta}^t s(\boldsymbol{y})\right\}}{\kappa(\boldsymbol{\theta}, \widehat{\boldsymbol{\phi}})} \widehat{f}_{\text{Laplace}}(\boldsymbol{y}|\widehat{\boldsymbol{\phi}}, \mu_\phi, \sigma^2_\phi). \tag{12}$$

The approximation in equation (12) is achieved using a Laplace approximation around the point $\widehat{\boldsymbol{\phi}}$. Details of this approximation are given in Section A of the appendix.
The Bayes factor of model m_2 against model m_1 is then defined as the ratio of (11) and (10), i.e.

$$\text{BF}_{21} = \frac{p(\boldsymbol{y}|m_2)}{p(\boldsymbol{y}|m_1)} = \frac{f(\boldsymbol{y}|\boldsymbol{\theta}, \mu_\phi, \sigma^2_\phi)}{f(\boldsymbol{y}|\boldsymbol{\theta}')} \cdot \frac{p(\boldsymbol{\theta})p(\mu_\phi)p(\sigma^2_\phi)}{p(\boldsymbol{\theta}')} \cdot \frac{p(\boldsymbol{\theta}'|\boldsymbol{y})}{p(\boldsymbol{\theta}, \mu_\phi, \sigma^2_\phi|\boldsymbol{y})}. \tag{13}$$

Applying the approximation from equation (12) to (13), and plugging in estimates for the posterior densities

$$p(\boldsymbol{\theta}'|\boldsymbol{y}) \approx \widehat{p}(\boldsymbol{\theta}'|\boldsymbol{y}) \quad \text{and} \quad p(\boldsymbol{\theta}, \mu_\phi, \sigma^2_\phi|\boldsymbol{y}) \approx \widehat{p}(\boldsymbol{\theta}, \mu_\phi, \sigma^2_\phi|\boldsymbol{y}) \tag{14}$$

leads to

$$\text{BF}_{21} \approx \frac{\exp\left\{\boldsymbol{\theta}^t s(\boldsymbol{y})\right\} \widehat{f}_{\text{Laplace}}(\boldsymbol{y}|\widehat{\boldsymbol{\phi}}, \mu_\phi, \sigma^2_\phi)}{\exp\left\{\boldsymbol{\theta}'^t s'(\boldsymbol{y})\right\}} \cdot \frac{\kappa(\boldsymbol{\theta}')}{\kappa(\boldsymbol{\theta}, \widehat{\boldsymbol{\phi}})} \cdot \frac{p(\boldsymbol{\theta})p(\mu_\phi)p(\sigma^2_\phi)}{p(\boldsymbol{\theta}')} \cdot \frac{\widehat{p}(\boldsymbol{\theta}'|\boldsymbol{y})}{\widehat{p}(\boldsymbol{\theta}, \mu_\phi, \sigma^2_\phi|\boldsymbol{y})}. \tag{15}$$

The ratio of the two normalizing constants $\kappa(\boldsymbol{\theta}')$ / $\kappa(\boldsymbol{\theta}, \widehat{\boldsymbol{\phi}})$ in (15) is estimated using a path sampling approach (Gelman and Meng, 1998), which is similarly used by Caimo and Friel (2013). Consider

$$\kappa(\boldsymbol{\theta}(g), \boldsymbol{\phi}(g)),$$

where

$$\boldsymbol{\theta}(g) = (1-g) \cdot \boldsymbol{\theta}' + g \cdot \begin{bmatrix} 0 \\ \boldsymbol{\theta} \end{bmatrix} \text{ and}$$
$$\boldsymbol{\phi}(g) = g \cdot \boldsymbol{\phi}$$

for $g \in [0, 1]$. So by construction

$$(\boldsymbol{\theta}(0), \boldsymbol{\phi}(0)) = (\boldsymbol{\theta}', \mathbf{0}) \quad \text{and} \quad (\boldsymbol{\theta}(1), \boldsymbol{\phi}(1)) = \left(\begin{bmatrix} 0 \\ \boldsymbol{\theta} \end{bmatrix}, \boldsymbol{\phi}\right).^2$$

Then thermodynamic integration (or so-called path sampling) can be used to estimate

$$\log\left\{\frac{\kappa(\boldsymbol{\theta}')}{\kappa(\boldsymbol{\theta}, \boldsymbol{\phi})}\right\} = \int_0^1 \mathbb{E}_{\boldsymbol{Y}|\boldsymbol{\theta}(g), \boldsymbol{\phi}(g)} \left[\left(\boldsymbol{\theta}' - \begin{bmatrix} 0 \\ \boldsymbol{\theta} \end{bmatrix}\right)^t s'(\boldsymbol{Y}) + (-\boldsymbol{\phi})^t t(\boldsymbol{Y})\right] \mathrm{d}g.$$

[2]Note that the additional 0 entry is necessary because the mixed effects model contains no parameter for the edges statistic. The edge effect is captured in the mean value μ_ϕ of the nodal random effects. The missing edges statistics is also the difference between $s'(\boldsymbol{y})$ and $s(\boldsymbol{y})$.

Consider discretising $g \in [0,1]$ as $(g_0 = 0, \ldots, g_i = \frac{i}{I}, \ldots, g_I = 1)$. Then we approximate

$$E_i := \mathbb{E}_{\boldsymbol{Y}|\boldsymbol{\theta}(g_i),\boldsymbol{\phi}(g_i)} \left[\left(\boldsymbol{\theta}' - \begin{bmatrix} 0 \\ \boldsymbol{\theta} \end{bmatrix} \right)^t s'(\boldsymbol{Y}) + (-\boldsymbol{\phi})^t t(\boldsymbol{Y}) \right]$$
$$\approx \frac{1}{N} \sum_{j=1}^{N} \left[\left(\boldsymbol{\theta}' - \begin{bmatrix} 0 \\ \boldsymbol{\theta} \end{bmatrix} \right)^t s'(\boldsymbol{y}^{(j)}) + (-\boldsymbol{\phi})^t t(\boldsymbol{y}^{(j)}) \right],$$

where the networks $\boldsymbol{y}^{(j)}$ are drawn from $f(\boldsymbol{y}|\boldsymbol{\theta}(g_i), \boldsymbol{\phi}(g_i))$, for $j = 1, \ldots, N$. Then we use a trapezoidal rule to numerically integrate

$$\log \left\{ \frac{\kappa(\boldsymbol{\theta}')}{\kappa(\boldsymbol{\theta}, \boldsymbol{\phi})} \right\} = \sum_{i=1}^{I-1} (g_{i+1} - g_i) \cdot \left(\frac{E_{i+1} + E_i}{2} \right).$$

The path sampling routine can easily be parallelised because the evaluations at the individual grid points of g do not depend on each other.

The Bayes factor in equation (15) is evaluated using the posterior mean values for the parameters $\boldsymbol{\theta}$, $\boldsymbol{\theta}'$, μ_ϕ and also for $\widehat{\boldsymbol{\phi}}$. For σ_ϕ^2 we plug in the mean of the logarithmized values and transform it back onto the scale of σ_ϕ^2, because the posterior density of σ_ϕ^2 is not symmetric.
For reasons of simplicity the posterior density estimates (14) are estimated assuming asymptotic normality, again using $\log(\sigma_\phi^2)$. For the data examples in the next section this assumption seems to be reasonable when looking at the plotted posterior density estimates. Furthermore, the individual contributions of the different components of the Bayes factor calculation suggest that at least in these cases the posterior density estimates play a minor part compared to the other components. If this assumption is violated this step in the algorithm can be changed.

4 Examples

4.1 Data Examples

4.1.1 Zachary's Karate Club Network

As a first data example we employ Zachary's karate club network (Zachary, 1977) which is a very well known data set often used in network analysis. The undirected 34 node network represents the friendships among members of a university karate club. Figure 2 shows a plot of this network graph. It is evident that there are only some nodes with a very high degree (no. 1, 33, and 34) while the majority of the remaining vertices has only two to four links. If there are no additional nodal attributes available, that might explain some differences between the actors, like for example status in the club (trainer, student, etc.), the assumption of vertex homogeneity in a standard ERGM appears to be at least questionable.

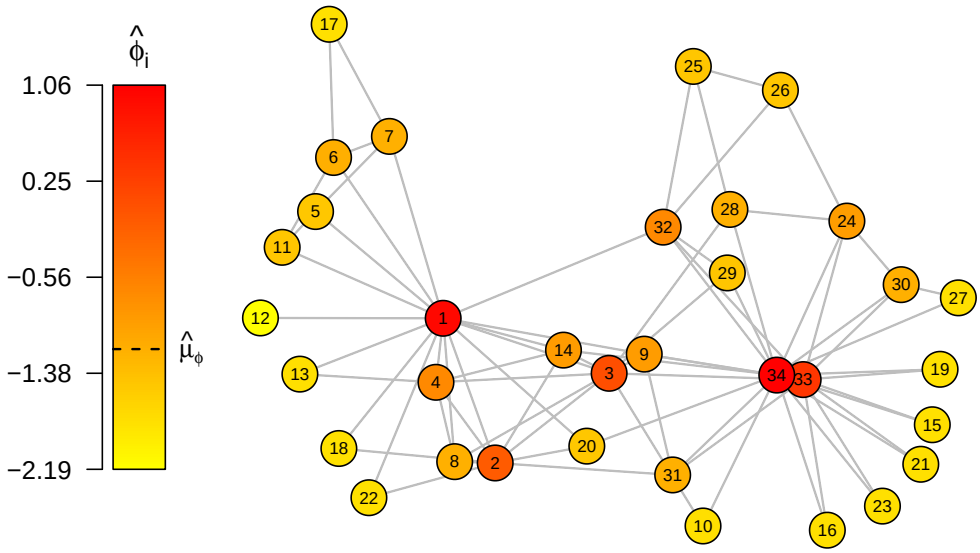

Figure 2 Zachary's Karate Club Graph. Vertices are coloured by their estimated nodal effect $\widehat{\phi}_i$ (posterior mean), $i = 1, \ldots, n$. Vertices with a high nodal effect are darker in orange/red.

We fitted two different models to the data: a standard ERGM with edges and triangles as sufficient statistics, and a model with nodal random effects and the triangle statistic. These two models are nested.
For the model fitting tasks we used the `Bergm` package (Caimo and Friel, 2014) and our extension of the `Bergm` routines, respectively. With 1,000 burn-in iterations, 30,000 main iterations, and 3,000 auxiliary iterations for the network simulation in each MCMC step, the computation of the fixed model took about two minutes on a 2.1 Ghz processor, the mixed model needed about one hour and forty minutes. Using 3,000 auxiliary iterations should be large enough because we have 561 possible ties in the network. Again, we refer to the results of Everitt (2012).
Table 1 shows the resulting posterior estimates for both models.

Figure 3 shows the results for the fixed model with edges and triangular effect only.

Figure 4 shows the results for the mixed model with nodal random and triangular effects for the karate club data.

Figure 5 shows estimates for the posterior densities for both models simultaneously.

What is evident from the estimated posterior densities is the difference for the triangular effect in both models in the upper right plot of Figure 5. When not accounting for nodal heterogeneity this effect is clearly positive compared to the mixed model where the posterior support clearly comprises zero. For the parameter θ_{edges} associated with the edges statistic and μ_ϕ in the mixed model there is no big difference between both models

Table 1 Model fitting results for the karate club data.

Model type	Parameter	Post. mean	Post. Sd.	Acceptance rate	Note
fixed	θ_{edges}	-2.32	0.16	0.43	
	$\theta_{\text{triangles}}$	0.54	0.11		
mixed	μ_ϕ	-1.17	0.22	0.26	
	σ^2_ϕ	1.05	0.58	0.54	*
	$\theta_{\text{triangles}}$	-0.04	0.21	0.09	

* For σ^2_ϕ the posterior mean is calculated based on the logarithmized values and than transformed back to the scale of σ^2_ϕ (this leads to the geometric mean) due to the non-symmetric posterior density in this case.

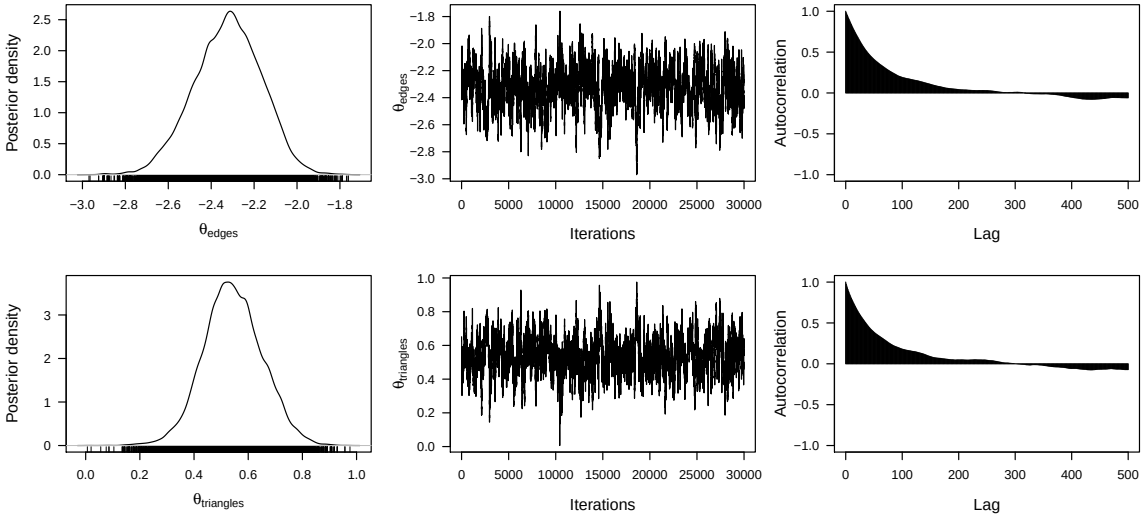

Figure 3 Posterior densities, trace plots, and autocorrelation for the fixed model with edges and triangular effect for the karate club data.

concerning the location of the posterior (when comparing θ_{edges} to $2 \cdot \mu_\phi$; note the different axis annotations).

The different effect of the triangular statistic in both models clearly illustrates the issue of model selection. After fitting the two competing models we computed a Bayes factor using the approach described in Section 3 to compare the model with nodal random effects to the one with structural effects only and tackle this issue.

The resulting estimated log Bayes factor is 453, which is huge. As explained in the previous section, there is some randomness involved in the procedure. Repeated calculation led to similarly huge values. This clearly indicates that the model with nodal random effects is preferable to the one without and is not surprising, because here a model with nodal heterogeneity appears much more realistic than one without.

Computing a single Bayes factor took about fourteen minutes using five 2.2 Ghz cores

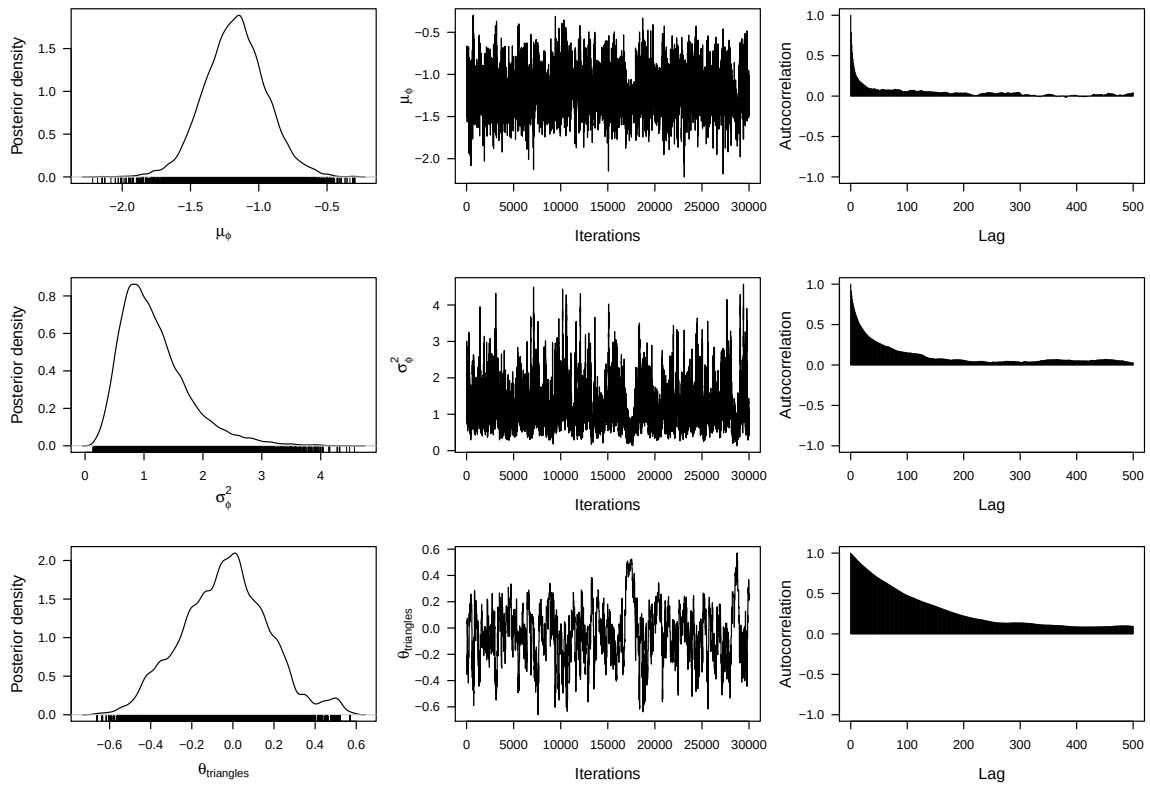

Figure 4 Posterior densities, trace plots, and autocorrelation for the mixed model with nodal random and triangular effects for the karate club data.

in parallel, with 10,000 iterations for the Laplace approximation, 1,000 grid points, 1,000 iterations at each point for the path sampling, and 3,000 iterations for each network simulation.

The vertices of the karate network in Figure 2 are coloured according to their estimated nodal effect $\widehat{\phi}_i$, $i = 1, \ldots, n$. As an estimate we use the corresponding posterior mean of each parameter ϕ_i. Vertices with a high value are darker in orange/red. By using such a colouring scheme we are able to visualise the variation in the nodal effects. In addition, we can identify important nodes in the network based on the estimated nodal effects.

4.1.2 European Parliament Members

The second data example consists of a network of members of the European parliament (MEP) in of the 6th legislative period. The complete network contains more than 900 vertices. We analyse a subset of the 32 members from the Netherlands. The induced subgraph is shown in Figure 6. A link between to MEPs exists if they have at least one committee membership in common. The data were provided by Paul W. Thurner (see Thurner et al., 2013). This data example illustrates our model selection procedure.

We fitted the same two nested models to the data as for the previous example: a standard

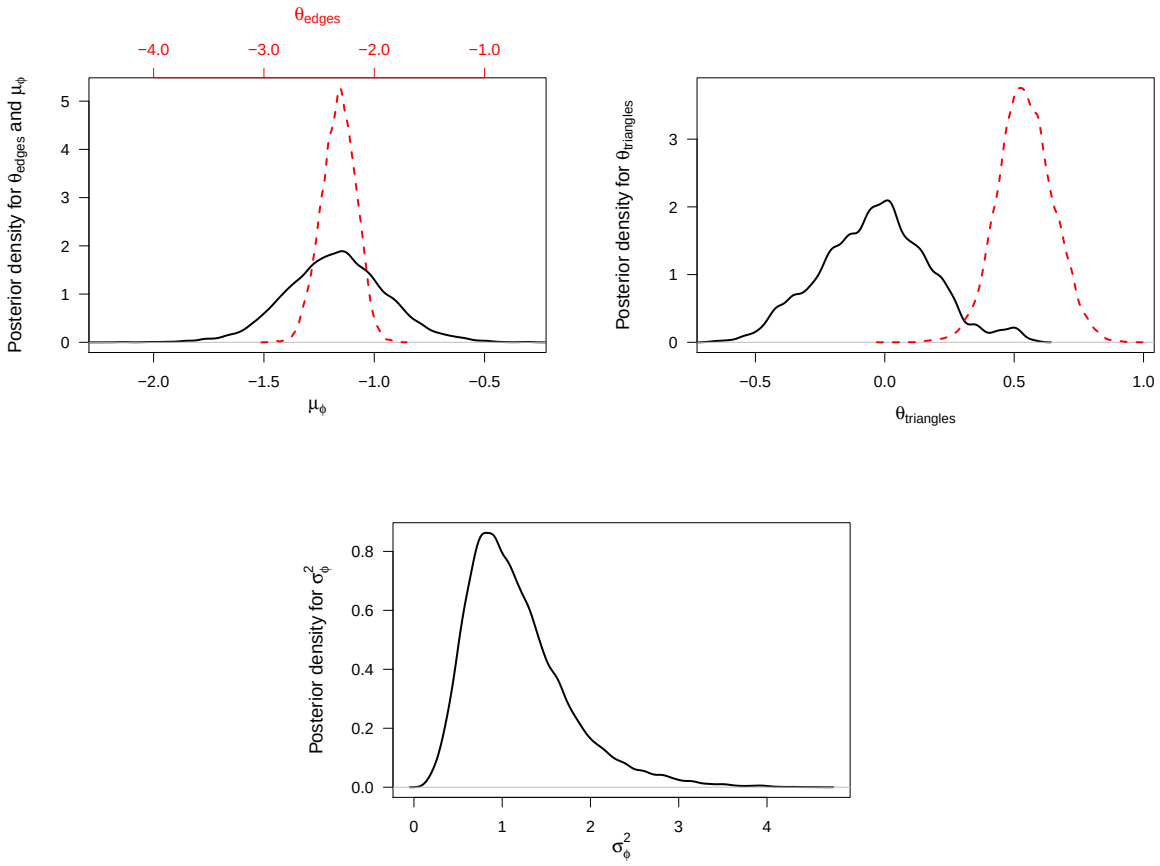

Figure 5 Posterior densities for the model with edges and triangular effect (red dashed lines) and nodal random and triangular effects (black solid lines) for the karate club data.

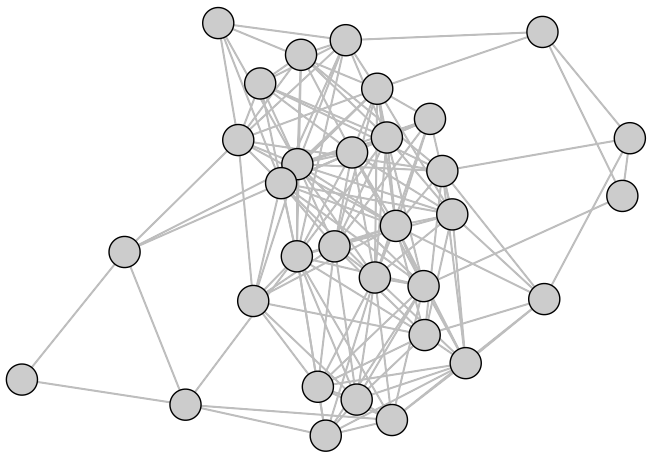

Figure 6 Network of Dutch members of the European parliament during the 6th legislative period. Two members are linked if they have at least one committee membership in common.

ERGM with edges and triangles as sufficient statistics, and a model with nodal random effects and the triangle statistic. The number of iterations was also equivalent.
Figures 7 and 8 show the results, which are also summarised in Table 2. For the mixed model we get a very low acceptance rate for the triangle effect and very high autocorrelations for the triangle effect and the mean parameter μ_ϕ. The later could possibly be solved by thinning out the chain.

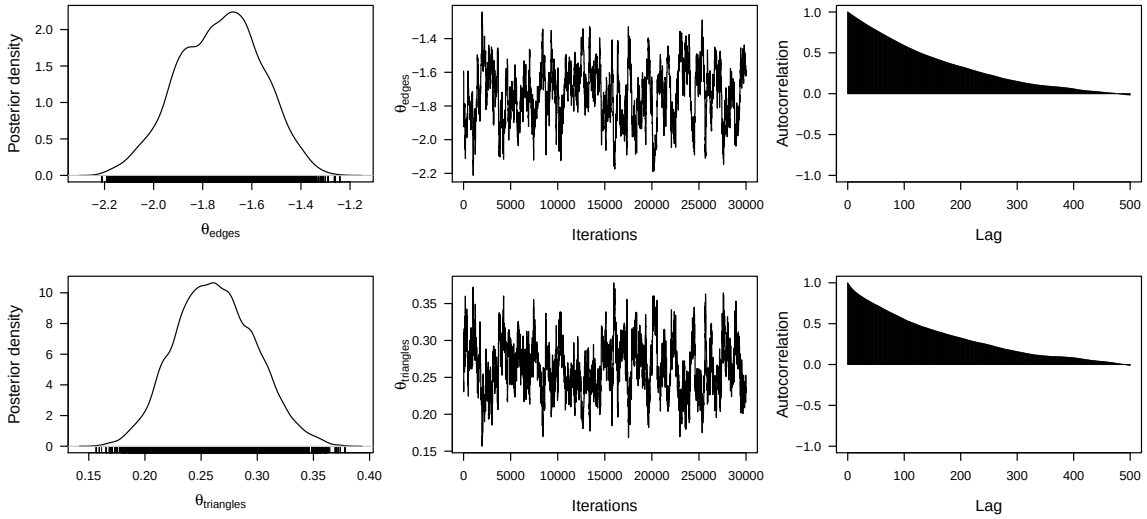

Figure 7 Posterior densities, trace plots, and autocorrelation for the fixed model with edges and triangular effect for the European parliament data.

Table 2 Model fitting results for the European parliament data.

Model type	Parameter	Post. mean	Post. Sd.	Acceptance rate	Note
fixed	θ_{edges}	-1.73	0.17	0.13	
	$\theta_{\text{triangles}}$	0.26	0.04		
mixed	μ_ϕ	-1.02	0.13	0.11	
	σ^2_ϕ	0.15	0.08	0.15	*
	$\theta_{\text{triangles}}$	0.29	0.05	0.02	

* For σ^2_ϕ the posterior mean is calculated based on the logarithmized values and than transformed back to the scale of σ^2_ϕ (this leads to the geometric mean) due to the non-symmetric posterior density in this case.

Nevertheless, the focus in this example is on model selection. The computed log Bayes factor is -13.9 and clearly indicates that the model without nodal random effects is preferable in this situation. Apparently, here we have a network dataset where there is no benefit in including nodal random effects into the model. This corresponds to the rather small estimate for the variance of the nodal random effects σ^2_ϕ. The resulting

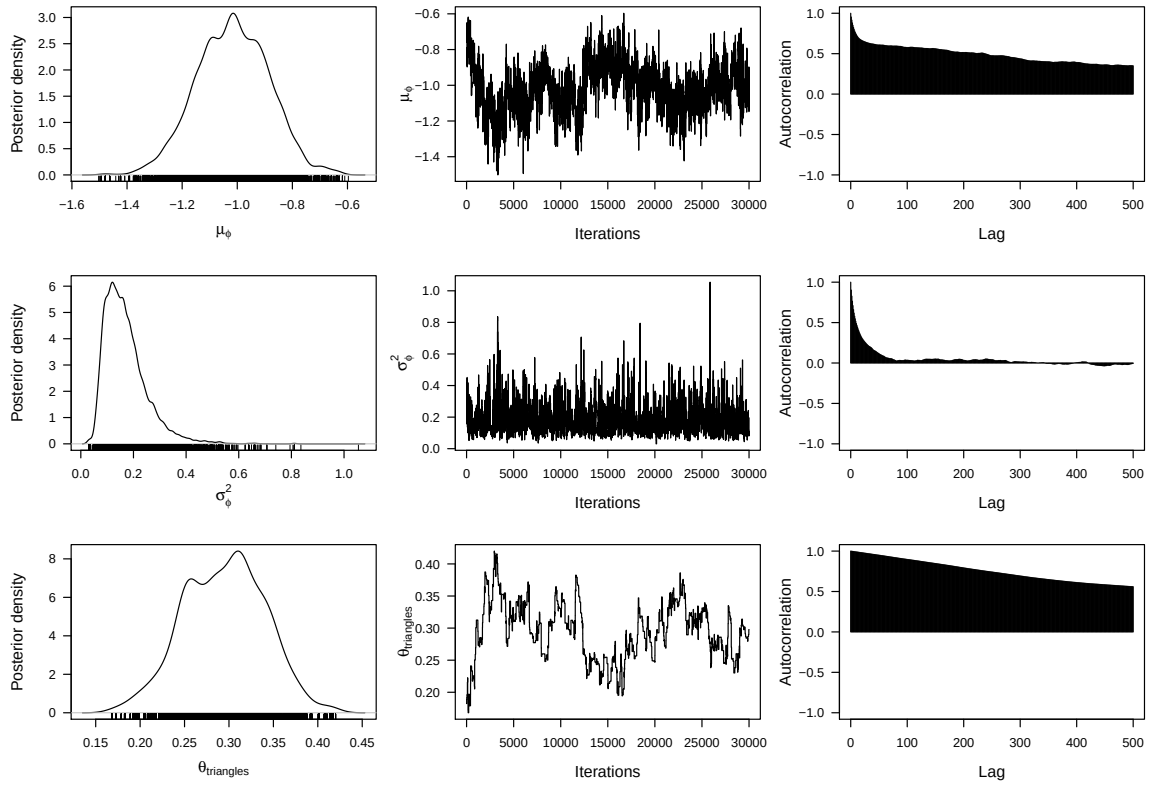

Figure 8 Posterior densities, trace plots, and autocorrelation for the mixed model with nodal random and triangular effects for the European parliament data.

Bayes factor shows that it is not the case that the model with more parameters is always selected. This can also be seen from the simulation results in the following subsection.

4.2 Simulation

For the simulation study we used the following components based on two very simple, but different model generating processes, a nodal random effects only situation, i.e. the p_2 model, and structural effects only situation, i.e. the classical ERGM. For each setting we generated networks with 40 vertices, using again the simulation routines from the `ergm` package (Hunter et al., 2008). The first model (A) was the one with nodal random effects only, i.e.

$$\text{logit}\left[\mathbb{P}\Big(Y_{ij}=1|Y_{kl},(k,l)\neq(i,j);\boldsymbol{\phi}\Big)\right]=\phi_i+\phi_j, \quad \text{with } \phi_i\sim N(\mu_\phi,\sigma^2_\phi), \quad \text{for } i=1,\ldots,n. \tag{16}$$

The parameter μ_ϕ was constantly set to $\mu_\phi=-1$, so that the resulting network graphs tend to be rather sparse. For σ^2_ϕ we used values between 0 and 1. Model (B) was the standard ERGM with edges and 2-star statistics, and no nodal random effects, i.e. $\boldsymbol{\theta}=(\theta_{\text{edges}},\theta_{\text{2-star}})^t$ and

$$\text{logit}\left[\mathbb{P}\Big(Y_{ij}=1|Y_{kl},(k,l)\neq(i,j);\boldsymbol{\theta}\Big)\right]=\theta_{\text{edges}}+\theta_{\text{2-star}}\cdot\left[\sum_{k\neq j}y_{ik}+\sum_{l\neq i}y_{jl}\right]. \tag{17}$$

The parameter θ_{edges} was constantly set to $\theta_{\text{edges}} = -2$. This is equivalent to model (A) in the sense that $2 \cdot \mu_\phi = \theta_{\text{edges}}$, because θ_{edges} is a parameter on a per link basis, μ_ϕ is on a per node basis and one needs two nodes to form a link. For $\theta_{\text{2-star}}$ we used values between 0 and 0.05. This value needs to be small, i.e. close to zero, because otherwise we only generate full or empty graphs if the value is negative, see also Schweinberger (2011). For each of the resulting parameter combinations in model (A) and model (B) we generated 50 networks.
For the chosen settings the resulting 40 node networks seem to be reasonable. We get an average network density between 0.11 and 0.30 for the different settings.
Note that setting $\sigma_\phi^2 = 0$ in model (A) and $\theta_{\text{2-star}} = 0$ in model (B) leads to a simple Bernoulli network, which can be seen as a null model.
Similarly to the karate data example we fitted two nested models to each of the simulated networks: a standard ERGM with edges and 2-stars as sufficient statistics, and a model with nodal random effects and the 2-star statistic. Again this step was followed by computing a Bayes factor to compare the model with nodal random effects to the one with structural effects only.

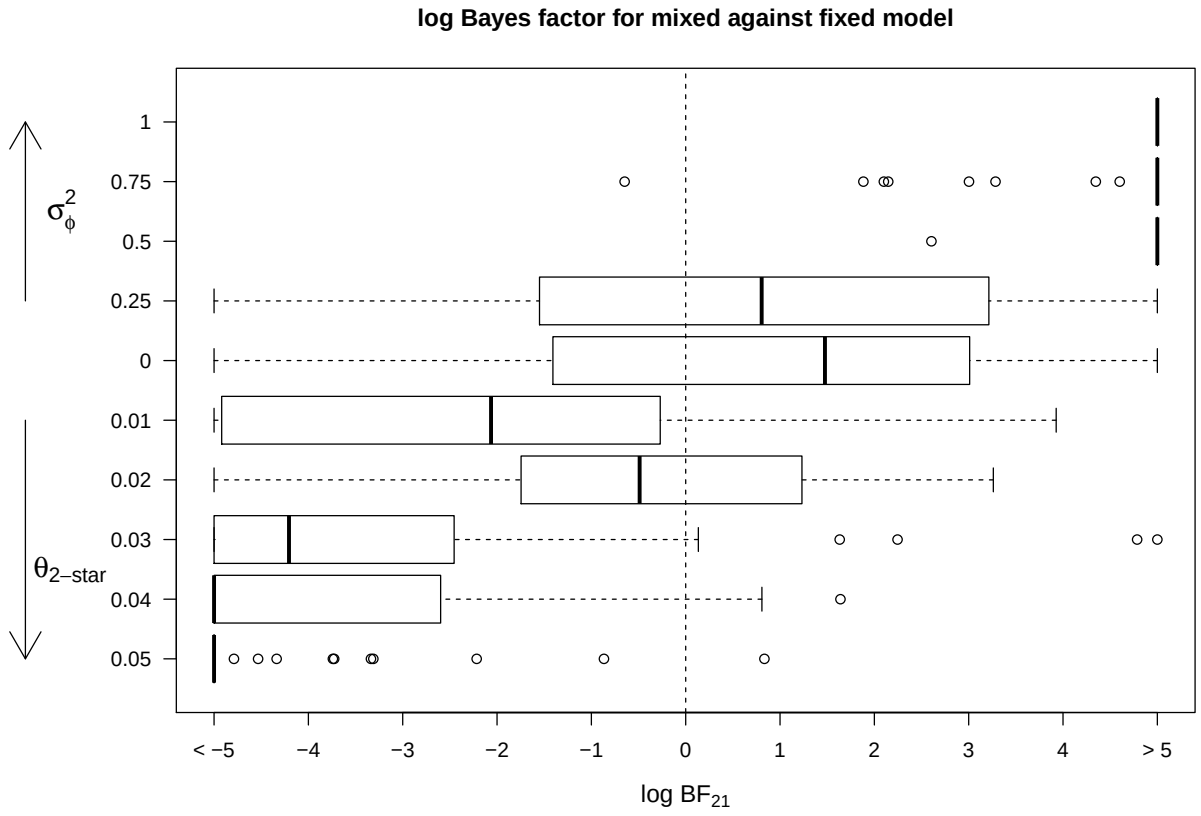

Figure 9 Resulting log Bayes factors for the mixed model against the fixed model for different simulation settings. The annotation on the y-axis shows which was the underlying true model, a model with nodal random effects only in the direction of σ_ϕ^2, and a model with edges and 2-stars in the direction of $\theta_{\text{2-star}}$.

Figure 9 shows boxplots of the resulting log Bayes factors for the different settings. For the plot the log Bayes factors were cut at values of -5 and 5 because some were really

small or really large. These cutting values were chosen following Kass and Raftery (1995). More detailed information, especially on the range of the simulation results is given Table 3. For the null model of a pure Bernoulli network the log Bayes factor can point in either one of the directions, the same is more or less true for only small deviations from this null model. The general impression is, that the more extreme the underlying setting becomes the sooner the log Bayes factor points into the correct direction.
Most importantly the results of the simulation show that our model selection works with respect to the size of the competing models. It is not the case that the model with more parameters, which is the model with nodal random effects, is always preferred.

Table 3 Resulting Bays factors (mixed model against fixed model) for the simulation from setting (A) a nodal random effects only situation, and setting (B) a classical ERGM with edges and 2-star statistics. Each setting was run 50 times, except for the Bernoulli setting, which had $2 \cdot 50$ *runs.*

Setting		average nw density	log Bayes factor for mixed against fixed model min	max	% <-5	% < 0	% > 0	% > 5
(A) random effects	$\sigma^2_\phi = 1$	0.23	13.03	137.64	0	0	100	100
	$\sigma^2_\phi = 0.75$	0.11	-0.65	498.53	0	2	98	84
	$\sigma^2_\phi = 0.5$	0.16	2.60	350.05	0	0	100	98
	$\sigma^2_\phi = 0.25$	0.15	-7.73	292.34	6	34	66	20
Bernoulli network $\sigma^2_\phi = \theta_{\text{2-star}} = 0$		0.13	-7.80	10.27	4	37	63	4
(B) fixed effects	$\theta_{\text{2-star}} = 0.01$	0.13	-14.56	3.93	24	76	24	0
	$\theta_{\text{2-star}} = 0.02$	0.14	-144.23	3.26	10	54	46	0
	$\theta_{\text{2-star}} = 0.03$	0.16	-25.24	51.34	44	88	12	2
	$\theta_{\text{2-star}} = 0.04$	0.20	-240.76	1.64	64	94	6	0
	$\theta_{\text{2-star}} = 0.05$	0.30	-218.07	0.83	80	98	2	0

Note: For setting (A) we set $\mu_\phi = -1$, and for setting (B) $\theta_{\text{edges}} = -2$, so that $\mu_\phi = 2 \cdot \theta_{\text{edges}}$.

5 Discussion and Summary

Statistical modelling of network data, with few exceptions, for example, Krivitsky et al. (2009), implicitly assumes that the local structure of the network is homogeneous. In particular, this implies that well studied phenomena, such as a small-world networks, Milgram (1967), Watts and Strogatz (1998), where shortest path lengths between two nodes in the network tend to be very small and scale-free networks, where few nodes have unusually high degree, are not appropriately modelled using the standard statistical modelling approaches. This is particularly true for Exponential Random Graph Models.

Here our extension of the Exponential Random Graph Model (ERGM) avoids the assumption of nodal homogeneity. By adding nodal random effects to the model we get a flexible tool to model heterogeneity in the network which is not captured in available (nodal) covariates otherwise. Using the Bayesian framework for ERGMs proposed by Caimo and Friel (2011) allows us to add this random effects extension to the model in an elegant and rather straightforward manner. Estimating Bayes factors enables us to handle the problem of model selection associated with this modelling task. The resulting estimates for the two data examples seem to be reasonable.
Furthermore, the small simulation study in the previous section suggests that in general the Bayes factor approach seems to work and even though a mixed model with nodal random effects has more parameters than its fixed equivalent it is not systematically preferred in the model selection.

We should note that the approach which we have introduced is computationally intensive. A promising avenue of research to address this issue is to explore approximations of the likelihood function using composite likelihoods, of which the pseudolikelihood approximation Frank and Strauss (1986) is an antecedent. We refer the reader to Varin et al. (2011) for a recent review of composite likelihoods. We are currently engaged in work in this direction.

Acknowledgements

We gratefully acknowledge the data provision by Paul W. Thurner.

The Insight Centre for Data Analytics is supported by Science Foundation Ireland under Grant Number SFI/12/RC/2289. Nial Friel's research was also supported by an Science Foundation Ireland grant: 12/IP/1424.

A Laplace approximation

The likelihood in the mixed effects model marginalized over the random effects $\boldsymbol{\phi}$ is

$$\begin{aligned} f(\boldsymbol{y}|\boldsymbol{\theta}, \mu_\phi, \sigma_\phi^2) &= \int \frac{\exp\left\{\boldsymbol{\theta}^t s(\boldsymbol{y}) + \boldsymbol{\phi}^t t(\boldsymbol{y})\right\}}{\kappa(\boldsymbol{\theta}, \boldsymbol{\phi})} \cdot p(\boldsymbol{\phi}|\mu_\phi, \sigma_\phi^2)\ \mathrm{d}\boldsymbol{\phi} \\ &= \int \frac{\exp\left\{\boldsymbol{\theta}^t s(\boldsymbol{y}) + \boldsymbol{\phi}^t t(\boldsymbol{y})\right\}}{\kappa(\boldsymbol{\theta}, \boldsymbol{\phi})} \cdot \frac{1}{(2\pi)^{\frac{n}{2}} \left|\sigma_\phi^2 I_n\right|^{\frac{1}{2}}} \cdot \\ &\quad \cdot \exp\left\{-\frac{1}{2\sigma_\phi^2}(\boldsymbol{\phi} - \mu_\phi \mathbb{1}_n)^t(\boldsymbol{\phi} - \mu_\phi \mathbb{1}_n)\right\} \mathrm{d}\boldsymbol{\phi} \\ &= \frac{\exp\left\{\boldsymbol{\theta}^t s(\boldsymbol{y})\right\}}{(2\pi\sigma_\phi^2)^{\frac{n}{2}}} \cdot \\ &\quad \cdot \int \exp\left\{\boldsymbol{\phi}^t t(\boldsymbol{y}) - \frac{1}{2\sigma_\phi^2}(\boldsymbol{\phi} - \mu_\phi \mathbb{1}_n)^t(\boldsymbol{\phi} - \mu_\phi \mathbb{1}_n) - \log(\kappa(\boldsymbol{\theta}, \boldsymbol{\phi}))\right\} \mathrm{d}\boldsymbol{\phi}. \end{aligned} \tag{18}$$

The integral in equation (18) is approximated around the point $\widehat{\boldsymbol{\phi}}$ using a Laplace type approximation

$$\int \exp\{h(\boldsymbol{\phi})\}\, \mathrm{d}\boldsymbol{\phi} \ \approx\ \exp\left\{h(\widehat{\boldsymbol{\phi}})\right\} (2\pi)^{\frac{n}{2}} \left|\Sigma\right|^{-\frac{1}{2}}, \tag{19}$$

where

$$h(\boldsymbol{\phi}) = \exp\left\{\boldsymbol{\phi}^t t(\boldsymbol{y}) - \frac{1}{2\sigma_\phi^2}(\boldsymbol{\phi} - \mu_\phi \mathbb{1}_n)^t(\boldsymbol{\phi} - \mu_\phi \mathbb{1}_n) - \log(\kappa(\boldsymbol{\theta}, \boldsymbol{\phi}))\right\}$$

and

$$\begin{aligned} \Sigma &= \frac{\partial^2 h(\widehat{\boldsymbol{\phi}})}{\partial\widehat{\boldsymbol{\phi}}\partial\widehat{\boldsymbol{\phi}}^t} \\ &= -\frac{1}{\sigma_\phi^2} I_n - \frac{\partial^2}{\partial\widehat{\boldsymbol{\phi}}\partial\widehat{\boldsymbol{\phi}}^t} \log(\kappa(\widehat{\boldsymbol{\theta}}, \boldsymbol{\phi})) \\ &= -\frac{1}{\sigma_\phi^2} I_n - \mathrm{Cov}(t(\boldsymbol{Y}), t(\boldsymbol{Y})^t|\widehat{\boldsymbol{\phi}}, \boldsymbol{\theta}). \end{aligned}$$

The matrix $\mathrm{Cov}(t(\boldsymbol{Y}), t(\boldsymbol{Y})^t|\widehat{\boldsymbol{\phi}}, \boldsymbol{\theta})$ denotes the covariance matrix of the vector of degree statistics $t(\boldsymbol{Y})$ and can be estimated via simulated networks using the parameters $\widehat{\boldsymbol{\phi}}$ and $\boldsymbol{\theta}$. These networks are drawn in the same way as the auxiliary networks needed for the exchange algorithm described in Section 2.

We assume that the posterior mode is close to the maximum likelihood estimator. The two are identical if the prior distributions are non-informative. This is not the case here, but we are assuming flat prior distributions and therefore the two should be reasonably close to each other. For reasons of simplicity, we use the posterior mean as value for $\widehat{\boldsymbol{\phi}}$.

Combining equation (18) with equation (19) yields

$$f(\boldsymbol{y}|\boldsymbol{\theta}, \mu_\phi, \sigma_\phi^2) \approx \frac{\exp\left\{\boldsymbol{\theta}^t s(\boldsymbol{y})\right\}}{\kappa(\boldsymbol{\theta}, \widehat{\boldsymbol{\phi}})} \widehat{f}_{\mathrm{Laplace}}(\boldsymbol{y}|\widehat{\boldsymbol{\phi}}, \mu_\phi, \sigma_\phi^2), \tag{20}$$

with

$$\widehat{f}_{\text{Laplace}}(\boldsymbol{y}|\widehat{\boldsymbol{\phi}}, \mu_\phi, \sigma_\phi^2) = \sigma_\phi^{-n} \exp\left\{\boldsymbol{\phi}^t t(\boldsymbol{y}) - \frac{1}{2\sigma_\phi^2}(\boldsymbol{\phi} - \mu_\phi \mathbb{1}_n)^t(\boldsymbol{\phi} - \mu_\phi \mathbb{1}_n)\right\} |\Sigma|^{-\frac{1}{2}} .$$

References

Caimo, A. and Friel, N. (2011). Bayesian inference for exponential random graph models. *Social Networks*, 33(1):41–55.

Caimo, A. and Friel, N. (2013). Bayesian model selection for exponential random graph models. *Social Networks*, 35(1):11–24.

Caimo, A. and Friel, N. (2014). Bergm: Bayesian exponential random graphs in R. *Journal of Statistical Software*, 61(2):1–25.

Duijn, M. A. J., Snijders, T. A. B., and Zijlstra, B. J. H. (2004). p_2: A random effects model with covariates for directed graphs. *Statistica Neerlandica*, 58(2):234–254.

Everitt, R. G. (2012). Bayesian parameter estimation for latent markov random fields and social networks. *Journal of Computational and Graphical Statistics*, 21(4):940–960.

Fellows, I. and Handcock, M. S. (2012). Exponential-family random network models. *arXiv preprint arXiv:1208.0121*.

Fienberg, S. E. (2012). A Brief History of Statistical Models for Network Analysis and Open Challenges. *Journal of Computational and Graphical Statistics*, 21(4):825–839.

Frank, O. and Strauss, D. (1986). Markov graphs. *Journal of the American Statistical Association*, 81(395):832–842.

Gelman, A. and Meng, X.-L. (1998). Simulating normalizing contants: from importance sampling to bridge sampling to path sampling. *Statistical Science*, 13(2):163–185.

Geyer, C. J. and Thompson, E. A. (1992). Constrained Monte Carlo maximum likelihood for dependent data. *Journal of the Royal Statistical Society. Series B (Methodological)*, 54(3):657–699.

Gill, P. S. and Swartz, T. B. (2004). Bayesian analysis of directed graphs data with application to social networks. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 53(2):249–260.

Goldenberg, A., Zheng, A. X., Fienberg, S. E., and Airoldi, E. M. (2010). A survey of statistical network models. *Foundations and Trends in Machine Learning*, 2(2):129–233.

Holland, P. W. and Leinhardt, S. (1981). An exponential family of probability distributions for directed graphs. *Journal of the American Statistical Association*, 76(373):33–50.

Hunter, D. R. and Handcock, M. S. (2006). Inference in curved exponential family models for networks. *Journal of Computational and Graphical Statistics*, 15(3):565–583.

Hunter, D. R., Handcock, M. S., Butts, C. T., Goodreau, S. M., and Morris, M. (2008). ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of Statistical Software*, 24(3):1–29.

Hunter, D. R., Krivitsky, P. N., and Schweinberger, M. (2012). Computational Statistical Methods for Social Network Analysis. *Journal of Computational and Graphical Statistics*, 21(4):856–882.

Kass, R. E. and Raftery, A. E. (1995). Bayes Factors. *Journal of the American Statistical Association*, 90(430):773–795.

Kolaczyk, E. D. (2009). *Statistical Anaysis of Network Models*. Springer, New York.

Krivitsky, P. N., Handcock, M. S., Raftery, A. E., and Hoff, P. D. (2009). Representing Degree Distributions, Clustering, and Homophily in Social Networks with Latent Cluster Random Effects Models. *Social Networks*, 31(3):204–213.

Milgram, S. (1967). The small world problem. *Psychology Today*, 2(1):60–67.

Murray, I., Ghahramani, Z., and MacKay, D. (2006). MCMC for doubly-intractable distributions. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI-06)*, Arlington, Virginia. AUAI Press.

R Core Team (2014). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Robins, G. L., Pattison, P., Kalish, Y., and Lusher, D. (2007a). An introduction to exponential random graph (p^*) models for social networks. *Social Networks*, 29(2):173–191.

Robins, G. L., Snijders, T. A. B., Wang, P., Handcock, M. S., and Pattison, P. (2007b). Recent developments in exponential random graph (p^*) models for social networks. *Social Networks*, 29(2):192–215.

Salter-Townshend, M., White, A., Gollini, I., and Murphy, T. B. (2012). Review of statistical network analysis: models, algorithms, and software. *Statistical Analysis and Data Mining*, 5(4):243–264.

Schweinberger, M. (2011). Instability, Sensitivity, and Degeneracy of Discrete Exponential Families. *Journal of the American Statistical Association*, 106(496):1361—-1370.

Snijders, T. A. B., Pattison, P. E., Robins, G. L., and Handcock, M. S. (2006). New specifications for exponential random graph models. *Sociological Methodology*, 36(1):99–153.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and van der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4):583–639.

Strauss, D. and Ikeda, M. (1990). Pseudolikelihood estimation for social networks. *Journal of the American Statistical Association*, 85(409):204–212.

Thurner, P. W., Kiel, M., and Schneider, M. (2013). Committee networks in the european parliament: Structure and impact on allocation of reports. Ms. LMU Munich.

Varin, C., Reid, N., and Firth, D. (2011). An overview of composite likelihood methods. *Statistica Sinica*, 21(1):5–42.

Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442.

Zachary, W. W. (1977). An information flow model for conflict and fission in small groups. *Journal of Anthropological Research*, 33(4):452–473.

Zijlstra, B. J. H., Duijn, M. A. J., and Snijders, T. A. B. (2006). The multilevel p_2 model: A random effects model for the analysis of multiple social networks. *Methodology: European Journal of Research Methods for the Behavioral and Social Sciences*, 2(1):42–47.